




MHS

Beyond Assessments

Ethics

At MHS, we are stewards of the personal information our customers provide. We use that data to understand individual behavior and develop effective interventions. As such, we believe that it is imperative that we have a well-articulated and easy-to-follow Trust Framework that governs how we interact with the data that we are entrusted with. The MHS Trust Framework has four pillars: Ethics, Stewardship, Transparency, and Accountability. This paper delves deeper into the Ethics pillar to provide more context around our ethics statement.

Ethics



When I started working on the overall Trust Framework for MHS, I saw ethics solely through my experience, training, and background in psychology and as a business leader. I would have defined it as “a set of guidelines that provide structure and rules to professions and professionals,” a moral code that somehow our training and education uniquely prepared us to abide by. After further research, it has become clear to me that I was oversimplifying a deep and complex topic.

Ethics, particularly the ethical use of data in an age of artificial intelligence (AI), algorithms, and prediction, has recently generated much discussion. Yet, debates about the ethical use of information and predictive results date back to the dawn of the computer age, with many schools of thought, each attempting to definitively provide a structure or framework that succinctly defines a very messy topic.

I grew up on a farm where things just had to get done. It’s in my DNA to work hard and to finish what I start. I was raised in an environment that was far too practical to enjoy lengthy debates that seemingly don’t have a right or wrong answer. You ate the food you grew, even if it included animals that you cared for. There was no right or wrong. There were just practicalities. As a result, until recently, philosophical topics have never been inherently interesting to me. Things changed as I assumed the CEO role for Multi-Health Systems (MHS) and led the effort to redefine a 40-year young success story for the digital age. It quickly became evident that I needed to explore further and explain what we meant by ethics in our Trust Framework. I had to embark on a journey deeper into territory that was not only unfamiliar but also, at times, uncomfortable.

To be clear, over my career, there have been many times when I have had to invoke my ethical principles to solve a business problem. However, until now,

I haven’t had to examine the foundations of my ethical principles or, more importantly, how they fit into the greater debate around ethics in a digital and data-driven global economy.

Ethics is the first of the four pillars of our MHS Trust Framework, with Data Stewardship, Transparency, and Accountability being the other three. Our statement on ethics reads:

“We commit to ensuring our use of data and the outcomes, predictions, prescriptions, and actions we take based on data will always be conducted ethically, ensuring rigorous attention is paid to ensure bias (conscious and unconscious) is removed from our products and services, always with the interest of the individual and their community as the guiding factor.”

Ethical Frameworks

Ethical theories can be organized into three main groupings

- Metaethics
- Normative ethics
- Applied ethics

Metaethics, as its name implies, looks at the big picture. Where do our ethical principles come from? What role do reason, the will of God, and other factors, play in establishing our ethical principles? Normative ethics takes a more practical approach and examines the moral standards that guide the establishment of right and wrong. In contrast, applied ethics concerns itself with examining specific controversies such as capital punishment, environmental concerns, and assisted suicide.

When attempting to resolve the issues raised in applied ethics, the frameworks provided by metaethics, and normative ethics are invoked as tools that can be used to provide answers or guide discussions. If we look at the world of artificial intelligence as an applied ethical issue, then we can use normative ethics and metaethics to guide us in resolving the inherent ethical dilemmas that arise when considering when, where, and how to apply AI techniques. For the purpose of this paper, I will leave metaethics aside since it offers little in terms of practical applications in day-to-day life and focus on normative ethics.

Ethics in an AI World

If artificial intelligence is an applied ethical issue and normative ethics map to the problem domain, the challenge is to identify which normative ethics principles we choose. Normative ethics can cover everything from rules and regulations provided by the government to professional standards and codes of conduct developed for specific industries to our own moral standards.

When it comes to the issue of ethics in modern complex topics such as AI and data-intensive machine learning techniques, it is only recently that legislative bodies have begun creating comprehensive rules and regulations to govern the use and application of these advanced technologies. As the media reports more on the outcomes that data and technology are creating, all levels of government will increasingly follow the lead of the European Union in drafting regulations on the use and transparency of data in decision-making.

With the ever-changing and sometimes contradictory regulations and legal guidelines, it may be beneficial to seek guidance from professional standards. MHS' solutions and services are primarily consumed by those in helping professions:

Psychologists, Business and Life coaches, Human Resources personnel, Probation and Parole workers, and Social Workers, for example. Many of these professionals have codes of conduct or rules that govern their actions. I don't know of any that specifically mention how to work with technologies such as AI. Still, I do know that they are all based on principles of "do no harm" with healthy doses of "respect your professional limits" and "client confidentiality" thrown in. They allow individuals a great deal of personal and professional judgment in applying the standards.

With these open-ended professional standards as guides, we are then left with our own individual standards, and by individual, I don't just mean us as single individuals but also as individual organizations. It is at this point that organizations, such as MHS, need to develop and implement ethical standards that govern how they operate in this increasingly complex environment.

Applied Ethics in Practice

How do we put this theory into practice? How does the above discussion on ethics help MHS to interpret or apply the ethics statement component of our trust framework, and how do our clients hold us accountable to our words?

Let's look at an example. Company A is hiring for a very technical role. They have a clearly defined set of competencies and job requirements. They can track the performance of existing people in the role and have performance metrics with solid objective measurements (remember, this is a hypothetical scenario). They have used historical performance data and the requirements for the new role to train an algorithm, allowing it to sort and prioritize applicants. To further expand the scenario, also consider that Company A has been called out in the industry for having very

poor representation of women and visible minorities in their technical employee pool, and as a result, has set diversity targets they are committed to meeting. In this situation, how do they ensure they get the best qualified person for the job while also meeting their diversity targets? What if two or more applicants are equally qualified, that is, they get the exact same result from the algorithm? Which person gets hired?

The reality is that most algorithms are not currently built or trained to handle decisions of this complexity nor to recognize the ethical component of the recommendation. In a hiring scenario, if two candidates are equally qualified according to the selection criteria, how does an algorithm prescribe which candidate to make an offer to? Further, algorithms also struggle with competing objectives, such as how to balance qualifications with diversity requirements.

In cases like this, it is essential that an algorithm does not make the final decision, operating with autonomy in the AI vernacular. Still, it makes a prescriptive recommendation that informs the decision of a subject matter expert. How can we encourage individuals to move beyond depending solely on algorithmic recommendations and instead incorporate the algorithm's input within a more comprehensive decision-making framework? After all, the comment that "because the computer says so" is often heard to justify action. What we seek is to retain decision-making by the "human in the loop," or the HIL as it is technically known.


AI systems offer several techniques to uphold decision-making at the human level. These approaches encompass the utilization of competing algorithms, which may yield different outcomes, and programming algorithms to incorporate probability factors, introducing an element of uncertainty. By employing these strategies, greater diversity and transparency can be achieved, enabling a thorough consideration of trade-offs to reach a conclusive decision. These techniques would not produce a single score but a range of results requiring a person to use judgment in decision-making. Organizational values and ethical principles can also be developed and applied, just as we are doing at MHS, to guide people away from over-relying on computer scores.

Retaining the Human in an Increasingly Automated World

Humans can understand, internalize and take actions or make decisions based on ethical principles. Whether we act on this capacity or not is a choice. AI systems, at least in their current form and thankfully absent from the capacity for artificial general intelligence (AGI), cannot make decisions on ethical principles. Given this, it is MHS' position to default, at least in areas that involve decisions with high-stakes outcomes for individuals, to the concept of Intelligent Assistance (IA) rather than Artificial Intelligence. We must maintain transparency into how a recommendation is arrived at and maintain a measured and re-creatable certainty index. When we reconceptualise AI as IA, we ensure that we retain a HIL to ensure so that the decision-making is not left entirely to algorithms; instead, we leverage algorithms to do the computational work, leaving humans to apply judgment. An example of how AI is re-conceptualized as IA is shown below (Figure 1: IA not AI).

Figure 1: IA not AI

AI in an IA capacity creates a great advantage for those who utilize it, but with great power comes even greater responsibility, which is why we adhere to the MHS Trust Framework in our actions, services, and products.



Have Questions? Get in touch with our team!